

Probabilistic Programming for Advancing Machine Learning

PROPOSERS' DAY

Arlington, VA

April 10, 2013





PPAML Proposers' Day Agenda

0900 – 1000 Check-In/Registration

1000 – 1005 Welcoming Remarks Dr. Kathleen Fisher, DARPA

1005 – 1020 Contracts Mr. Mark Jones, DARPA

1020 – 1030 Security Ms. Natalie Young, DARPA

1030 – 1130 PPAML Program Dr. Kathleen Fisher, DARPA

1130 – 1300 LUNCH BREAK

1300 – 1400 Individual Company Presentations

1400 – 1500 Government Response to Questions Dr. Kathleen Fisher, DARPA



Opening Remarks

Kathleen Fisher, DARPA Program Manager



Logistics

- DARPA-BAA 13-31
 - Posted on FedBizOpps website (<http://www.fedbizopps.gov>) and Grants.gov website (<http://www.grants.gov>)
 - Posting Date: April 1, 2013
 - Proposal Due Date: May 16, 2013 at Noon ET
- Procedure for Questions/Answer
 - Questions can be submitted until 1230
 - Questions will be answered during Q&A session in the afternoon
- Program website ([http://www.darpa.mil/Our_Work/I2O/Programs/Probabilistic_Programming_for_Advanced_Machine_Learning_\(PPAML\).aspx](http://www.darpa.mil/Our_Work/I2O/Programs/Probabilistic_Programming_for_Advanced_Machine_Learning_(PPAML).aspx))
 - Copy of presentations
 - Video recording
 - Frequently Asked Questions (FAQs)



Contracts

Mark Jones, DARPA Contracts Management Office



PPAML Proposers' Day

BAA PROCESS OVERVIEW

- Solicitation will be released utilizing BAA procedures in accordance with FAR 35.016
- The BAA is posted (as will any needed amendments) on FEDBIZOPPS at www.fbo.gov and Grants.gov at www.grants.gov.
- BAA allows for a variety of technical solutions.
- Proposal evaluations will be accomplished through a scientific review using the evaluation criteria stated in the BAA.
- The BAA contains one closing time/date – 12:00 noon ET, May 16, 2013.
- BAA covers all info needed to propose. Following the proposal preparation instructions assists the evaluation team to clearly understand what is being proposed and supports a timely conclusion to the selection process



PPAML Proposers' Day

ELIGIBILITY

- All interested/qualified sources may respond subject to the parameters outlined in BAA
- Foreign participants/resources may participate to the extent allowed by applicable Security Regulations, Export Control Laws, Non-Disclosure Agreements, etc.
- FFRDCs and Government entities are subject to applicable direct competition limitations and cannot propose to this BAA in any capacity, unless they clearly demonstrate the work is NOT otherwise available from the private sector AND provide written documentation citing the specific statutory authority establishing eligibility to propose to Government solicitations and (for FFRDCs) written authorization from the sponsoring agency
- Procurement Integrity: Potential Conflicts of Interest – Identify and discuss mitigation – failure to do so will result in proposal rejection without technical evaluation or further consideration for award



PPAML Proposers' Day

POTENTIAL AWARD INSTRUMENTS

FAR Based Procurement Contracts, Cooperative Agreements or Other Transactions. No Grants

TECHNICAL AREA AWARDS

- 4 Technical Areas (TA): TA 1, 2, 3, and 4
- Proposals should only cover one TA, with the exception that TA 2 and 4 can be submitted together in a single proposal.
- Proposers may propose to more than one TA, but entities selected to perform any task within TA 1 will not be selected in any capacity to perform in the other TA
- There are no restrictions on the number of non TA 1 technical areas that can be awarded to a single proposer
- Anticipate multiple awards under TA 2-4, only a single award under TA 1



PPAML Proposers' Day

PROPOSAL PREPARATION INFORMATION

- Consists of two volumes – Technical (with required Appendix A and optional Appendix B) and Cost
- Volume I - Technical and Management
 - Volume I has a 16 page limitation for the technical portion (Executive Summary, Goals and Impact, and Technical Plan) and a 10 page limitation for the management portion (any remaining Volume I sections). If submitting a combination TA 2/4 proposal, the technical portion limitation is increased to a total of 23 pages. The evaluation team will not review any submitted pages that exceed the Volume I limit.
 - Volume I includes a mandatory Appendix A and an optional Appendix B, neither appendix counts towards Volume I's page limit total.
- Volume II – Cost – No page limitation.
- BAA describes the necessary information to address is each volume –
 - Make sure to include every section identified
 - If section does not apply – put “None” (e.g., Animal Use – None, OCI - None)
 - Include a working spreadsheet as part of your Cost Volume submission
 - Notionally, TA 1 has a different starting schedule than other TAs
 - Remember: Appendix A is mandatory



PPAML Proposers' Day

PROPOSAL PREP - TECHNICAL DATA RIGHTS

- Government desires, at a minimum, **Government Purpose Rights** for any proposed noncommercial software (including source code), software documentation, hardware designs and documentation, and technical data.
- Data Rights Assertions – Assert rights to all technical data & computer software generated, developed, and/or delivered to which the Government will receive **less than Unlimited Rights**. This information may be assessed during evaluations.
 - Provide and justify basis of assertions that apply to the Prime and any Subs. A prescribed format will be included in the BAA. Break out these assertions in a separate table (if possible) to be included as an attachment to a resultant contract or agreement.
 - Explain how the Government will be able to reach its program goals (including transition) within the proprietary model offered; and
 - Provide possible nonproprietary alternatives in any areas that might present transition difficulties or increased risk or cost to the Government under the proposed proprietary solution. NOTE: Offerors expecting to use, but not to deliver, open source tools or other materials in implementing their approach may be required to indemnify the Government against any legal liability arising from such use.



PPAML Proposers' Day

ITEMS TO NOTE

- Understand and be compliant with the System for Award Management (SAM), Electronic and Information Technology compliance, Employment Eligibility Verification (E-verify), Reporting Executive Compensation and First-Tier Subcontract Awards and Updates of Information Regarding Responsibility Matters (FAPIIS)
- Awardees will be required to use i-Edison, T-FIMS and Wide Area Workflow (WAWF)
- Subcontracting Issues
 - **NON SMALL BUSINESSES:** Subcontracting Plans required for FAR based contracts with subcontracting possibilities expected to exceed \$650,000
 - Subcontractor cost - Proposals must include, at a minimum, a non-proprietary, subcontractor proposal for EACH subcontractor
 - If utilizing FFRDC, Government entity, or a foreign owned firm as a subcontractor, submit their required eligibility information as applicable
 - Providing Accelerated Payment to Small Business Subcontractors Clause Deviation



PPAML Proposers' Day

- Proposals must be valid for a minimum of 120 days
- If a prospective proposer believes a conflict of interest exists or may exist (whether organizational or otherwise) or has a question on what constitutes a conflict, the proposer should promptly raise the issue with DARPA by sending the proposer's contact information and a summary of the potential conflict to the BAA mailbox before preparing a proposal and mitigation plan.
- Document files must be in Portable Document Format (.pdf, ISO 32000-1), OpenDocument (.odx, ISO/IEC 26300:2006), .doc, .docx, .xls, or .xlsx formats.
- Submissions must be written in English.



PPAML Proposers' Day

PROPOSAL SUBMISSION

- DARPA anticipates submitted proposals will be UNCLASSIFIED. Classified proposals will not be accepted.
- Follow the procedures outlined in the BAA for Procurement Contract / Other Transaction submissions vs. Cooperative Agreement submissions
- If submitting multiple proposals through the DARPA web-based upload system, a unique user ID and password must be created for each submission. If submitting through the Grants.gov website, following their procedures for submitting multiple proposals.
- DO NOT submit proposals by any electronic means – This includes classified email/fax machine submissions.
- DO NOT wait until the last minute to submit proposals – the submission deadline is strictly enforced and late submissions may not be evaluated.



PPAML Proposers' Day

EVALUATION / AWARD

- No common Statement of Work - Proposal evaluated on individual merit and relevance as it relates to the stated research goals/objectives rather than against each other.
- Evaluation Criteria are, in descending order of importance: (a) Overall Scientific and Technical Merit; (b) Potential Contribution and Relevance to the DARPA Mission; and (c) Cost Realism.
- Evaluation Process is a scientific/technical review - Reviews conducted by panels of experts that may include contracted Government SETAs bound by strict non disclosure agreements.
- Government reserves the right to select for award all, some, or none of the proposals received, to award portions of a proposal, and to award with or without discussions. Contracts may be either classified or unclassified.
- No portion of this announcement will be set aside for Historically Black Colleges and Universities (HBCUs), Small Businesses, Small Disadvantaged Businesses and Minority Institutions (MIs) and no preferences apply.



PPAML Proposers' Day

COMMUNICATION

- Prior to Receipt of Proposals – No restrictions, however Gov't (PM/PCO) shall not dictate solutions or transfer technology. Unclassified FAQs will be periodically posted to this BAA's DARPA Web page. Classified FAQs (if applicable) will be faxed via a classified fax machine to all registered/eligible parties.
- After Receipt of Proposals – Prior to Selection: Government (PM/PCO) may communicate with offerors in order to understand the meaning of some aspect of the proposal that is not clear or to obtain confirmation or substantiation of a proposed approach, solution, or cost estimate.
- After Selection/Prior to Award: Government (PCO) may clarify aspects of the proposal and/or may conduct negotiations. Government (PM/COR/PCO) may clarify the Statement of Work or, in cases where only portions of the proposal are accepted, may discuss reductions to the scope to match the selected effort.
- Informal feedback for non selected proposals may be provided once the selection(s) are made.

Only a duly authorized Contracting Officer may obligate the Government



Security

Natalie Young, Program Security Representative



Meeting Classification Level

UNCLASSIFIED DISCUSSIONS ONLY



Security Requirements

The overall classification of Programming for Advancing Machine Learning (PPAML) is UNCLASSIFIED

Proprietary and Classified Information

- Proprietary Information: Proposers are responsible for clearly identifying proprietary information. Submissions containing proprietary information must have a proprietary information cover page with each page containing proprietary information clearly marked.
- Classified Information: Because the program emphasizes the idea of creating and leveraging open source technology, classified submissions WILL NOT be accepted under this solicitation.



Performer Approval to Release Information

Controlled Unclassified Information (CUI) on Non-DoD Information Systems

- CUI is unclassified information that does not meet the standard for National Security Classification but is pertinent to the nation interests of the United States or to the important interest of entities outside the Federal Government and under law or policy requires
 - (1) Protection from unauthorized disclosure
 - (2) Special Handling Safeguards
 - (3) Prescribed Limits on Exchange or Dissemination

Performer Responsibilities:

- In accordance with DARPA Instruction 65 (DI 65), material considered to be contracted fundamental research does not require public release review. To determine if your submission requires a review, please review your contract with DARPA.



DARPA Security Points of Contact

PROGRAM SECURITY REPRESENTATIVE

- Natalie Young
 - 703-812-1980
 - natalie.young.ctr@darpa.mil

PROGRAM SECURITY OFFICER

- Gregory Woosley
- (571) 218-4538
- gregory.woosley@darpa.mil



Probabilistic Programming for Advancing Machine Learning (PPAML)

Kathleen Fisher, DARPA Program Manager



Machine Learning is Ubiquitous and Very Useful

ISR



DARPA Grand Challenge
Fully autonomous ground vehicles competition

Nuclear Test Ban Treaty Compliance: Deduce set of seismic events given detections and misdetections on a network of stations

Image Search/Activity Detection: Find and identify objects and actions in video

Object Tracking: Follow vehicles as they move through a city and are recorded in multiple video streams (DARPA CZTS)

Patterns of Life: Process wide area aerial surveillance data and associated tracks to infer location types and object dynamics

Bird Migration Patterns: Model spatio-temporal distribution of birds (by species); involves large-scale sensor integration

DARPA LAGR: Vision-based robot navigation

Google Glasses: Perform searches based on images taken by user cell phone cameras

Natural Language Processing



© Apple

Siri
Voice recognition and Natural Language Processing (NLP)

Watson: Computer system capable of answering questions posed in natural language

Topic Models: Statistical model for discovering the abstract "topics" that occur in a collection of documents

Distributed Topic Models: Asynchronous distributed topic discovery

Citation Analysis: Given citations, extract author, title, and venue strings and identify co-reference citations

Entity Resolution: Discovering entities that are mentioned in a large corpus of news articles

NLP Sequence Tagging: Tagging parts of speech and recognizing named entities in newspaper articles

Predictive Analytics



© Netflix

Netflix Challenge Predict user ratings for films based on previous ratings

Microsoft Matchbox: Match players based on their gaming skill set

Predictive Database: Understand information based on causal relationships in data

Bing Image Search: Search for images on the web by selecting text in word document

Amazon Recommendation Engine: Recommend items based on consumer data

Cyber and Other



ORNL's Attack Variant Detector: Discover compromised systems

Yahoo's Bayesian Spam Filtering: Self-adapting system based on word probabilities

Cyber Genome Lineage: Reverse engineer malware samples to find shared "genetic" features between different malware samples

Gene Sequencing: Determine order of nucleotides in a DNA molecule

Disclaimer: Images of specific products are used for illustration only. Use of these images does not imply endorsement of inherent technical vulnerabilities.



Explaining the Magic by Example: Channel State Estimation

Phenomenon:

Channel properties between cell tower and cell phones, each with >1 antenna (MIMO)

Question: Given received signal, what was sent?

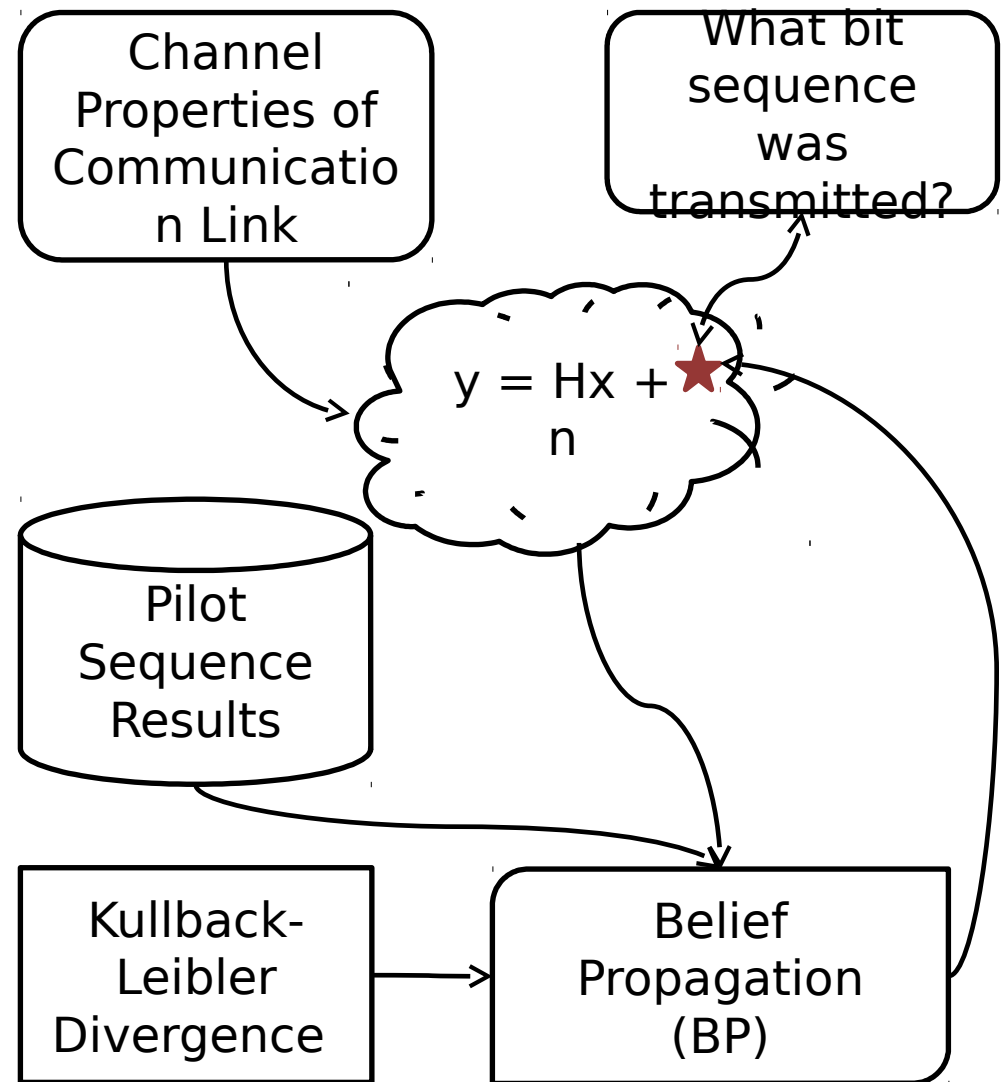
Training Data:

Received signals given pilot (known) transmissions

ProtoModels:

Channel matrix H plus vector n of Gaussian noise.

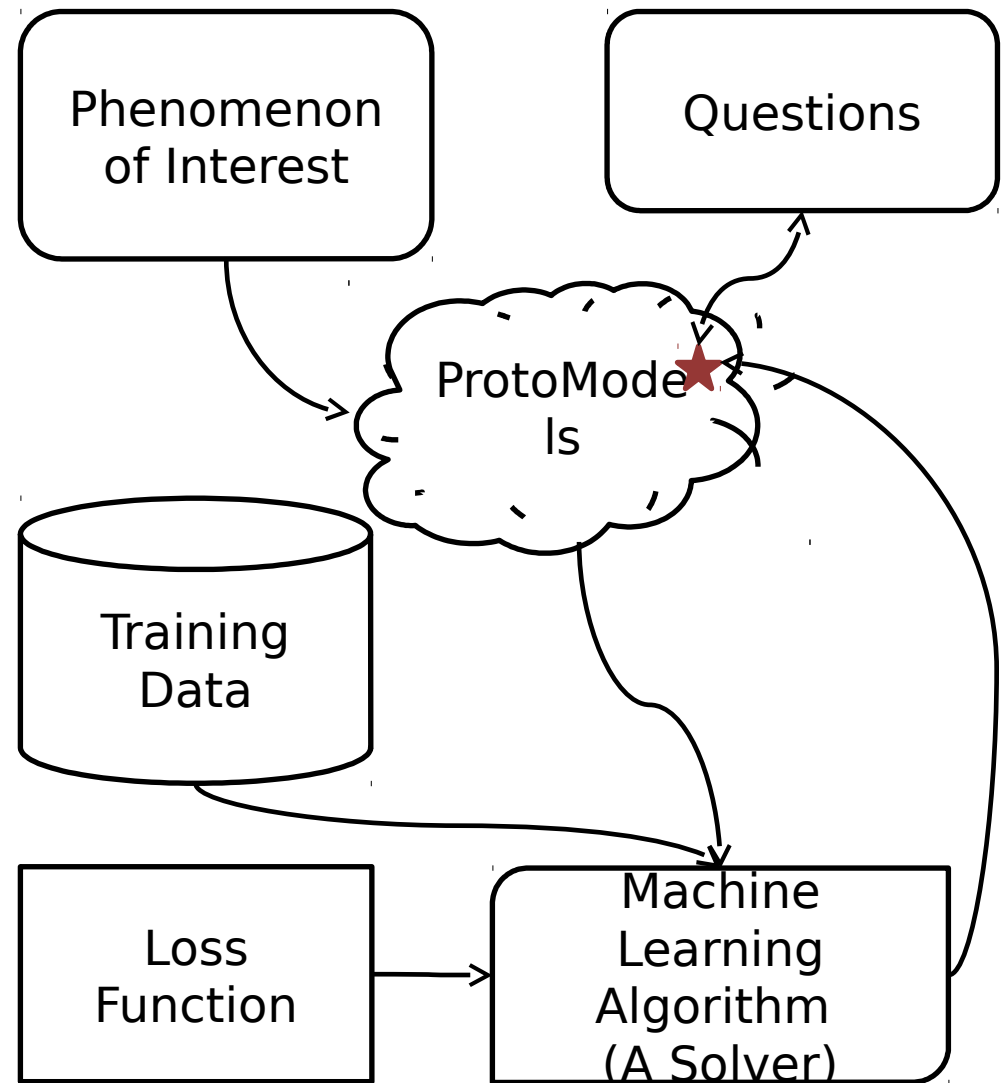
Selected **Model** approximates noise and channel matrix





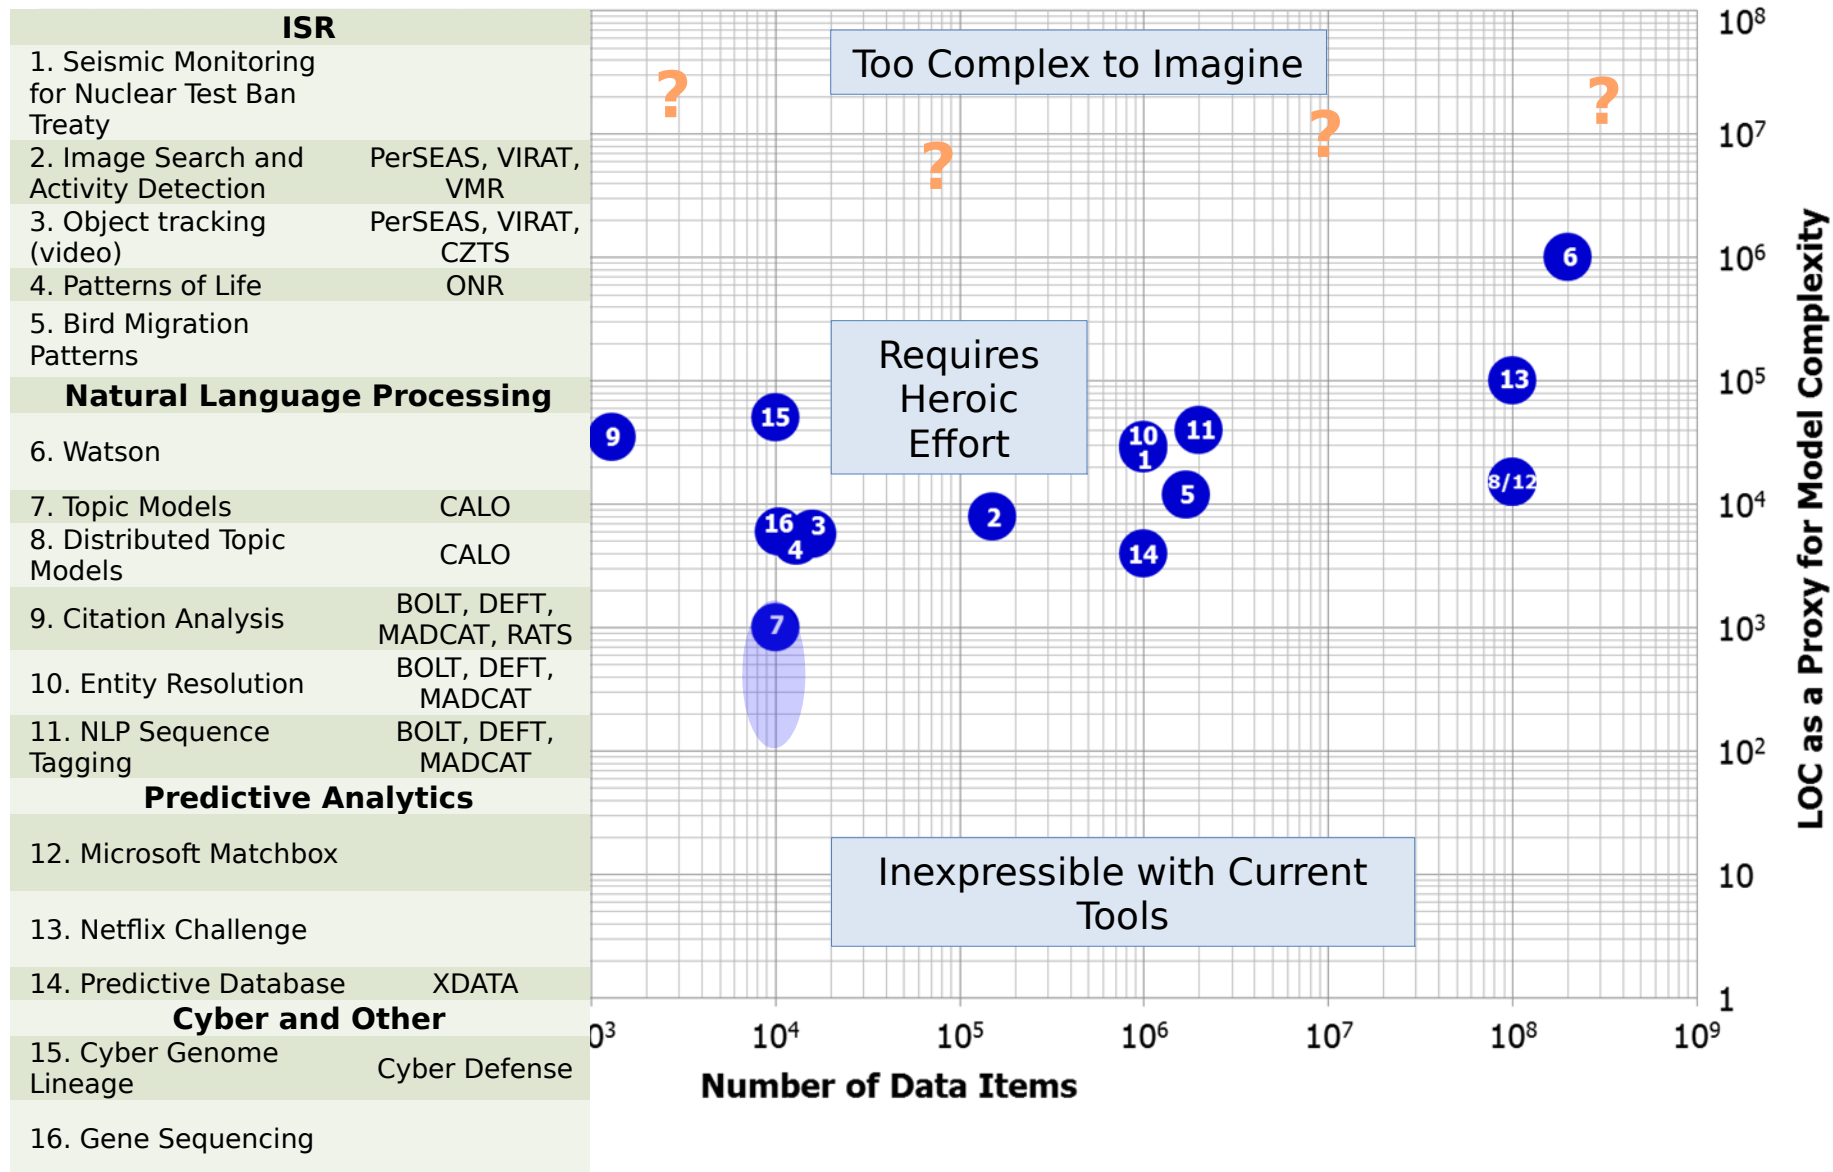
Why Is It Hard?

- Brittleness of implementations & lack of reusable tools. *PHY decoder: \$200M/standard; Infer.NET: max 20% on model*
- High level of required expertise *10K solvers, 100s of grad student hours per model*
- Painfully slow & unpredictable solvers *Massive data sets, complex algorithms, tricky coding for graph traversal and numeric stability*
- Challenges constructing models *Limited modeling vocabulary; models entwined with solvers*





We're Missing a Tool to Write These Applications





The Missing Tool (Explained by Example)

Model:

```
mem strength person = gaussian 100 10
lazy person = flip 0.1
pulling person = if lazy person then (strength person) / 2
                    else strength person
total-pulling team = sum (map pulling team)
winner team1 team2 = greater (total-pulling team1)
                      (total-pulling team2)
```



Source: Noah Goodman, POPL (2013)

Query:

strength Bob

Facts:

```
[Bob, Mark] = winner [Bob, Mark] [Tom, Sam]
[Bob, Fred] = winner [Bob, Fred] [Jon, Jim]
```

System will calculate probability distribution for Bob's strength given known facts

The user describes the model at a high level. An inference engine analyzes the program, query, data, and available hardware resources to produce best solution





Richer Example: Seismic Monitoring for Nuclear Test Ban Compliance

Goal: Deduce a bulletin listing seismic events with time, location, depth, and magnitude

Given: All the seismic detections and misdetections observed by a network of stations

Comparison:

- Existing UN System: > 28K LOC in C; took several years to build; cost ~\$100M
- NET-VISA: 25 LOC in C; took less than 1 hour*; cost \$400K



Source: Stuart Russell, NIPS Probabilistic Programming Workshop (2012)

Probabilistic Programming could make ML applications *better* and *easier to build*



The Probabilistic Programming Revolution

Traditional Programming *Probabilistic Programming*

Code models capture how the data was generated using random variables to represent uncertainty

Libraries contain common model components: Markov chains, deep belief networks, etc.

PPL provides probabilistic primitives & traditional PL constructs so users can express model, queries, and data

Inference engine analyzes probabilistic program and chooses appropriate solver(s) for available hardware

Hardware can include multi-core, GPU, cloud-based resources, GraphLab, UPSIDE/Analog Logic results, etc.

High-level programming languages facilitate building complex systems
Probabilistic programming languages facilitate building rich ML applications



The Promise of Probabilistic Programming Languages

- **Shorter:** Reduce LOC by 100x for machine learning applications
 - Seismic Monitoring: 28K LOC in C vs. 25 LOC in BLOG
 - Microsoft MatchBox: 15K LOC in C# vs. 300 LOC in Fun
- **Faster:** Reduce development time by 100x
 - Seismic Monitoring: Several years vs. 1 hour
 - Microsoft TrueSkill: Six months for competent developer vs. 2 hours with Infer.Net
 - Enable quick exploration of many models
- **More Informative:** Develop models that are 10x more sophisticated
 - Enable surprising, new applications
 - Incorporate rich domain-knowledge
 - Produce more accurate answers
 - Require less data
 - Increase robustness with respect to noise
 - Increase ability to cope with contradiction
- **With less expertise:** Enable 100x more programmers
 - Separate the model (the program) from the solvers (the compiler), enabling *domain experts* without machine learning PhDs to write

Sources:

- Bayesian Data Analysis, Gelman, 2003
- Pattern Recognition and Machine Learning, Bishop, 2007
- Science, Tanenbaum et al, 2011

Probabilistic Programming could empower domain experts *and* ML experts




Research Challenges

- Design probabilistic programming languages and end-user tools
 - Expressiveness vs. performance
 - Usability by end-users who are not machine learning experts
 - Profiling and debugging tools
- Create model/query/prior-data analyses to determine best solver or combination of solvers for given problem
- Build an inferencing infrastructure to efficiently “solve” high-level models
 - **Improve solver performance** by leveraging research from PL community
 - Define an API so new solvers can be slotted in
 - Develop new solvers
 - Leverage strengths of different kinds of hardware, cloud, G5, ...
- Develop a broader community
 - Enable all modelers to connect to all solvers to leverage everyone’s efforts
 - Develop model libraries for popular model building patterns

PL/ML collaborations are already bearing fruit:
Example shows 100x performance improvement

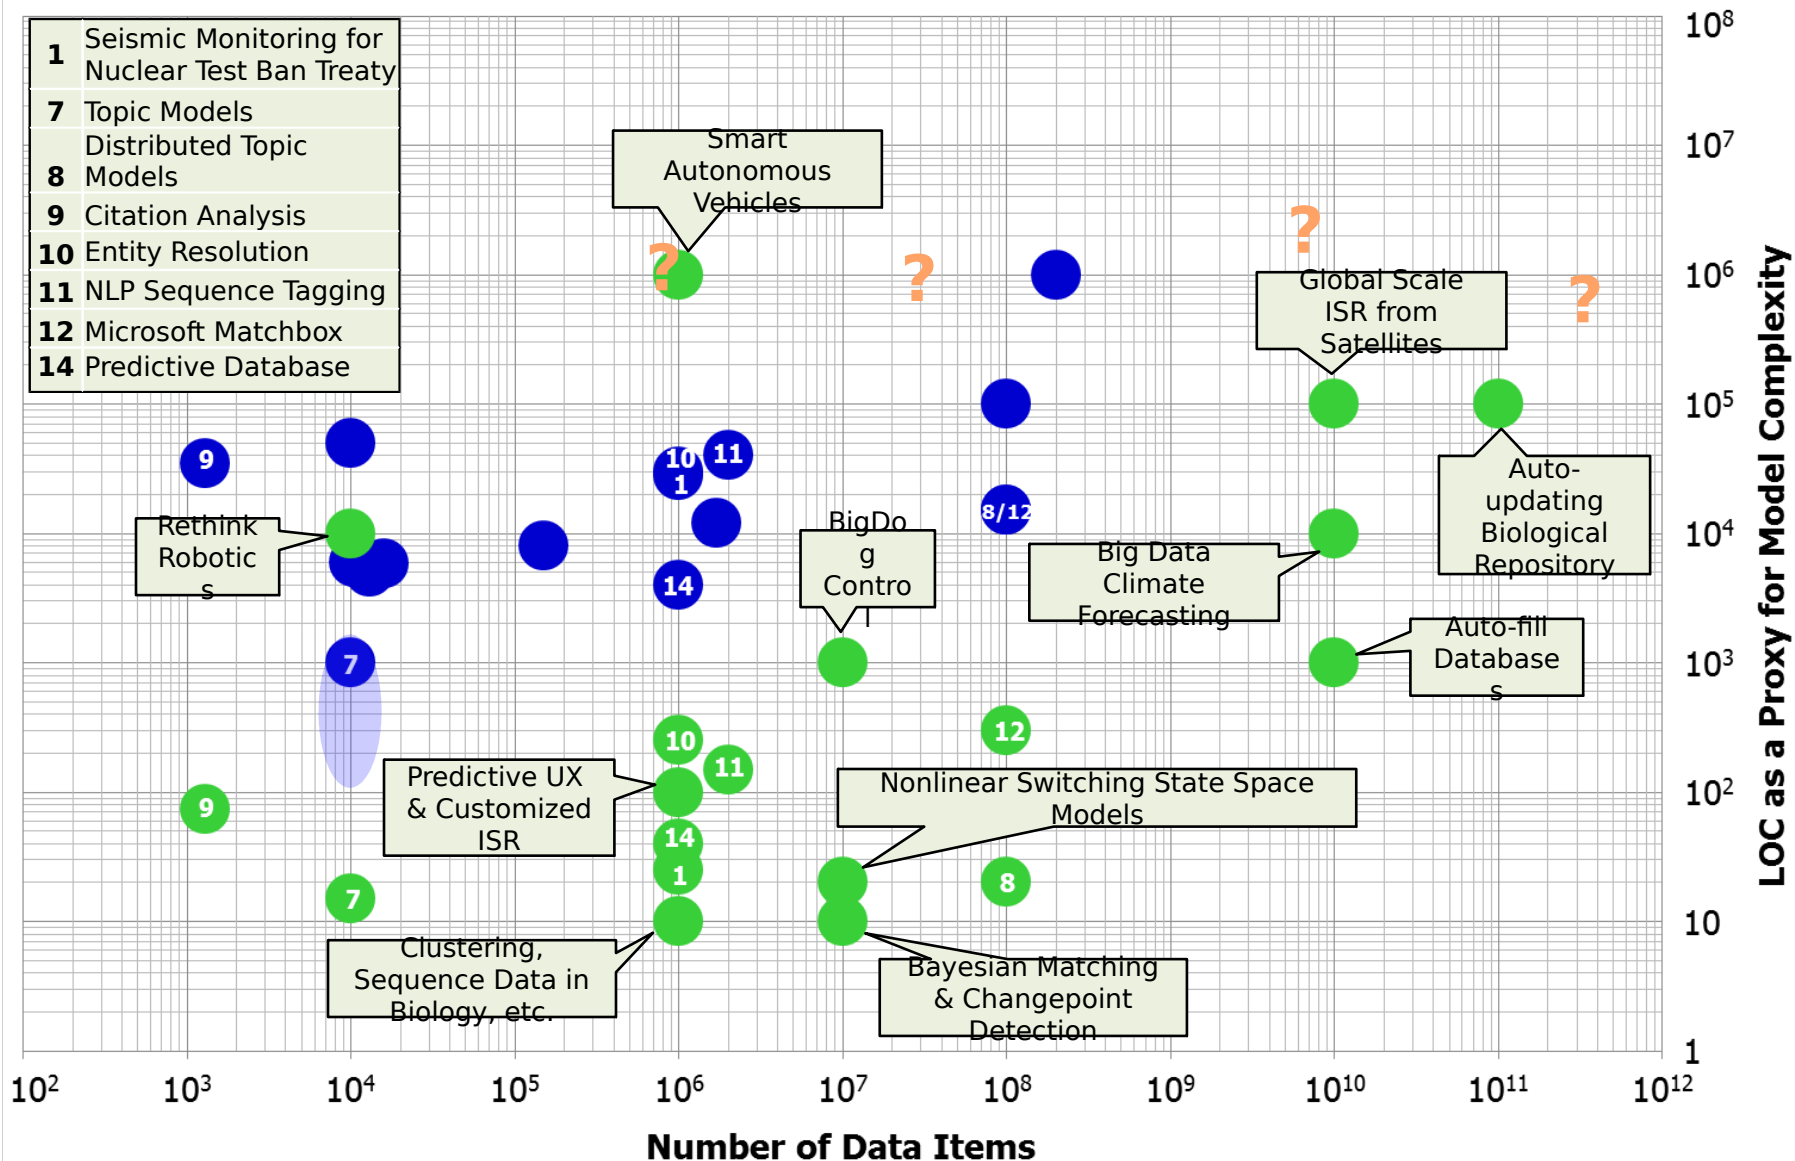
Performance comparison on different model sizes



Source: Noah Goodman, PLDI (2013)

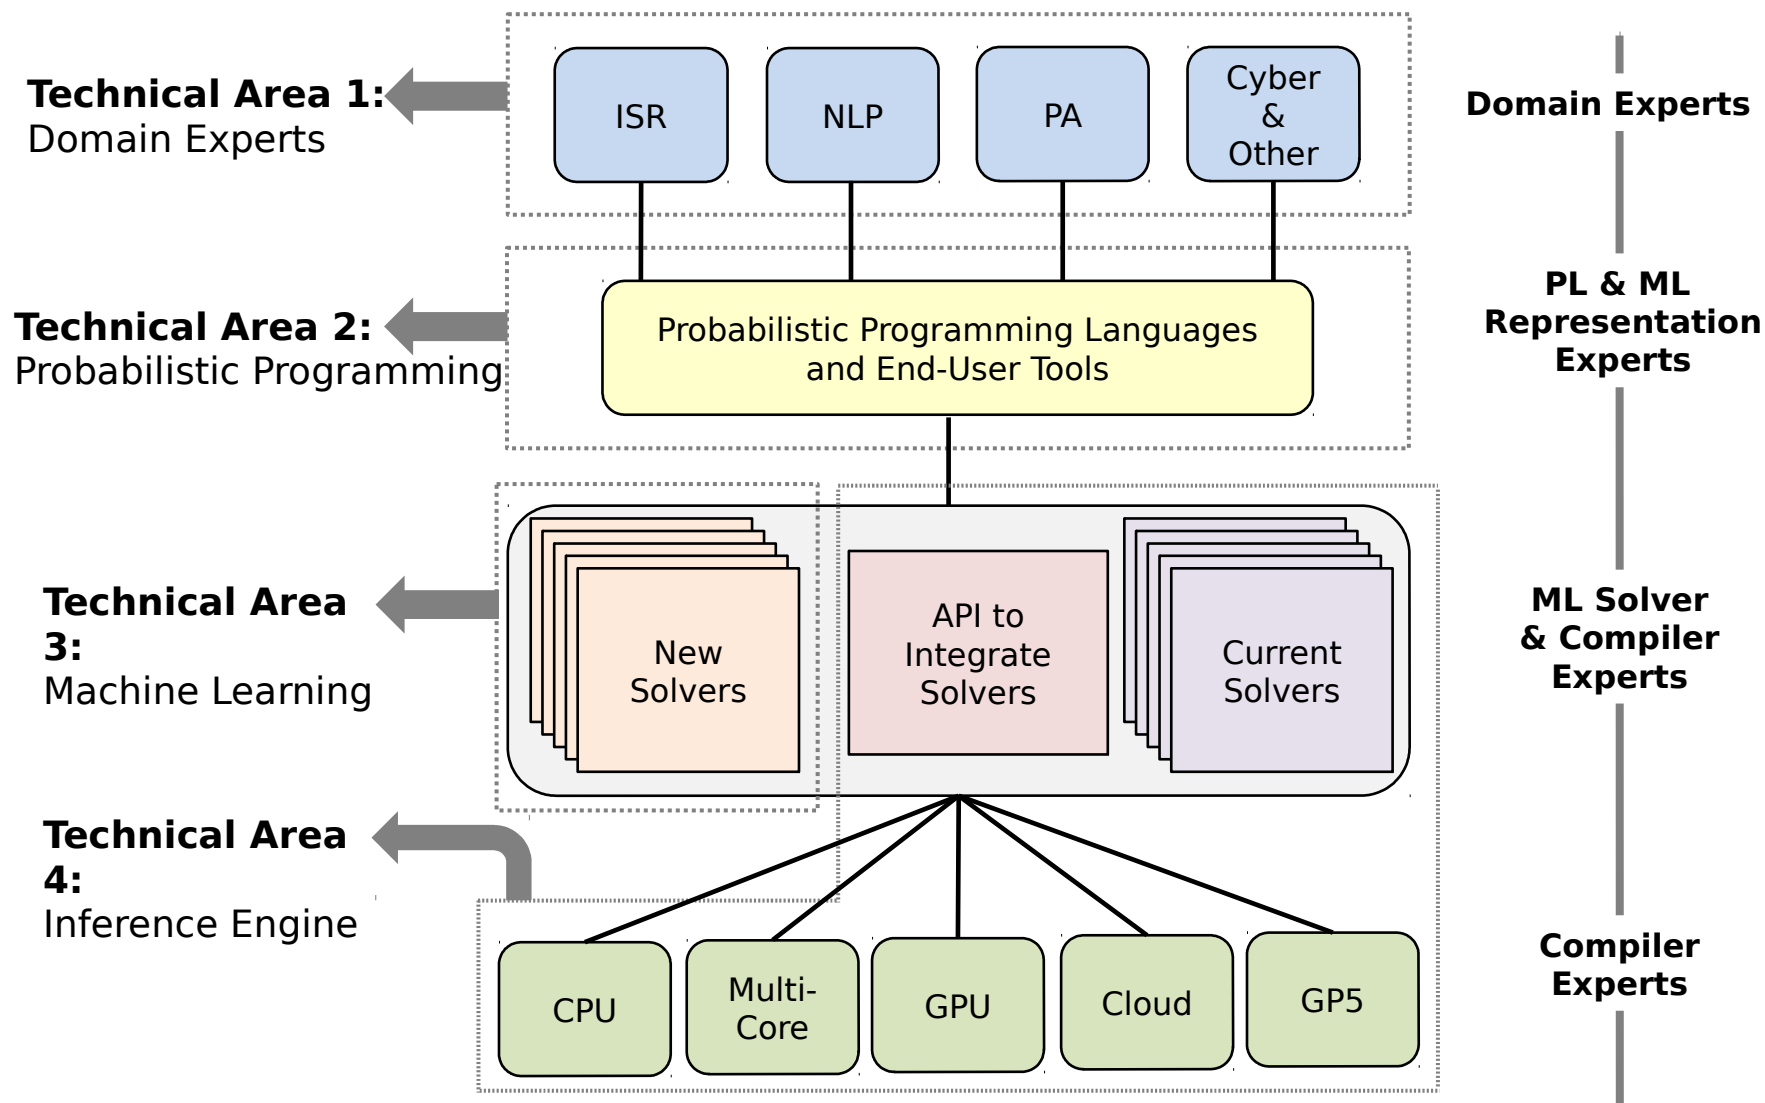


If This Technology Is Wildly Successful...






Proposed Program Structure





Notional Proposed Program Schedule, Products & Transition Plan



Products

- Probabilistic languages and end-user tools
- Extensible, high-performance compilers
- Library of reusable models
- Supporting community

Transition Plan

- Would include government customers from ISR, NLP, Predictive Analytics, Cyber and other domains in annual "Summer Schools" and in challenge problem selection



Programmatics

- Three planned phases: Phase I (22 months); Phases II & III (12 months each)
- Four Technical Areas (TAs)
 - TA1 - Domain Experts
 - TA2 - Probabilistic Programming
 - TA3 - Machine Learning
 - TA4 - Inference Engine
- Anticipate one award for TA1 and multiple awards in each of TA2-4
 - If selected for TA1, cannot be selected for any portion of the other TAs
- Performers in TA2-4 will be grouped into one or more design teams
 - Each design team led by a TA2 performer
 - Produce a working end-to-end Probabilistic Programming System (PPS)
 - Identify and produce working version of Team Challenge Problem (CP)
 - Performers may participate on more than one team
 - Teams will not be competitively evaluated; no anticipated down-selection
- Strong interaction among all performers is critical to program success
 - Associate Contractor Agreement (ACA)



Technical Area 1 – Domain Experts

- Develop Challenge Problems (CPs) from a diverse set of domains for TA2-4 performers
 - 3 at program kickoff; 1 new every 6 months
 - Grow in complexity; increase data amount
- Evaluate performance of each PPS on each Challenge Problem
 - Quality of the solution and run-time performance
 - Present findings at each PI meeting
- Organize and run annual Summer Schools
 - Identify attendees and associated ML problems
 - Select and contract with appropriate venues
 - Provide necessary computing resources
 - Provide access to relevant data sets
- Evaluate the effectiveness of each PPS at each Summer School
 - Usability; quality of solution; time required to produce solution, and performance of



- Make available experts in the domain of the CPs available for consultation with [TA2-4 performers](#)
- Make available to [TA2-4 performers](#) any infrastructure necessary for developing and evaluating the effectiveness of the solution
- Consult with the Government Team and consider feedback from [TA2-4 performers](#) and Summer School participants in selecting CPs beyond those defined at program kick-off
- Foster information sharing and collaboration amongst [all program performers](#) via online



Technical Area 2 – Probabilistic Programming

- Design and build “front-end” of a PPS
 - Design suitable probabilistic programming language in which users with a range of skills can concisely express powerful generative models
 - Develop supporting infrastructure to help end users understand the performance, correctness and accuracy consequences of their modeling decisions
 - Profilers
 - Debuggers
 - Model verification/checking tools, etc.
 - Develop language pragmas to enable experts to tune solver performance
- Demonstrate PPS on Team CP at each PI mtg
- Provide training and support to Summer School participants
- Demonstrate techniques on CPs and at Summer Schools



- Lead a PPS design team
 - One or more TA4 performer
 - Zero or more TA3 performers
 - Develop Team CP
- With TA4, define appropriate interfaces and common IRs to enable integration of technologies from TA2-4
- Deliver PPS to TA1 performers
- Provide technical support to TA1 performers during evaluation of the PPS on the program-wide CPs
- Incorporate novel algorithms, representations and analyses discovered by the TA3 performers



Technical Area 3 – Machine Learning

- Perform basic ML research that supports PPAML program goals
- Work on fundamental ML research problems identified by other PPAML performers in the course of the program
- Some potential research challenges:
 - Develop the theory of probabilistic programming
 - Discover new inference algorithms that are more efficient, accurate, predictable or generalizable
 - Discover novel representations that support more efficient, accurate, predictable or generalizable inferences
 - Develop inference algorithms that work over streaming data or have better scaling properties
 - Develop techniques for assessing model fitness for a particular data set



- Participate in one or more [PPS design teams](#), led by TA2
- Provide [support to TA2 and TA4 teammates](#) to integrate the novel algorithms, representations or analyses discovered into the developed Probabilistic Programming Systems
- Use the CPs developed by the [TA1 performers](#) and the Team CP to evaluate and drive research
- Work with [performers from all TAs](#) to identify critical research challenges and issues



Technical Area 4 – Inference Engine

- Design and build the “back-end” of a PPS
 - Inputs: generative models written in a probabilistic programming language; queries; and prior data
 - Output: efficient implementation with predictable performance
- Develop analyses to determine which solver or set of solvers is most appropriate for given input
- Improve efficiency of solvers, potentially by applying optimization techniques from PL/compiler community
- Compile inference engines to a range of different hardware targets: multi-core machines, GPUs, cloud-based clusters, novel hardware, etc.
 - Note: Designing new hardware is out of scope



- Participate in one or more **PPS design teams, led by TA2**
- Provide inference engine for the PPS developed by their **design team(s)**
- Incorporate novel algorithms, representations, and analyses discovered by the **TA3 performers**
- Use the CPs developed by the **TA1 performers** and the Team CP to evaluate and drive research
- Work with **performers from all TAs** to identify critical research challenges and issues